Towards Explaining Exchange Traded Funds' Impact on Market Volatility Using an Agent-based Model

Megan Shearer[1,3]
shearerj@umich.edu

David Byrd[2,3]
db@gatech.edu

Tucker Hybinette Balch[3]
tucker.balch@jpmchase.com

[1]University of Michigan, Ann Arbor
[2]Georgia Institute of Technology
[3]J.P. Morgan AI Research

Abstract

Exchange traded funds (ETF), which are baskets of securities that trade on the stock market, have become a popular investment tool, but potentially increase the price co-movement of their underlying securities. We study index-based ETFs, because underlying securities trade on the same market as their corresponding ETF, which creates ample opportunities for traders to arbitrage between price deviations in the ETF and the market index. Considering the increase in market volatility in 2018 and 2019 in the US stock market, we examine if an ETF spreads mini flash crashes, or small volatility events, from one of its underlying symbols to another. We address this question in a simulated environment with an ETF and two symbols which compose the ETF's portfolio. We explore two market environments, one with an ETF and arbitrage agents, and one without an ETF and arbitrage agents. We find that the presence of an ETF and arbitrage agents increases underlying symbols' overall volatility, and when one symbol experiences a mini flash crash, the other symbol does experience a momentary price change in the opposite direction. This supports the idea that ETFs potentially increase market-wide volatility by spreading volatility events through their portfolios. To further this work, we hope to utilize learning to implement smarter arbitrage agents and smarter background agents who might avoid adverse selection by not trading during mini flash crashes.

1 Introduction

Exchange traded funds (ETF) have become a popular investment vehicle to easily and efficiently gain more market exposure. An *ETF* is a portfolio of securities that trades on the stock market. The *underlying securities* in an ETF's portfolio can be any traded entity, such as stocks, bonds, or commodities. Any investor who can access the stock market themselves or through a broker, can also trade ETFs, making them more accessible to average investors than other funds [Antoniewicz and Heinrichs, 2014]. ETFs also provide opportunities for lower trading costs, lower risk, and lower volatility [Golub et al., 2013].

This work focuses on *index-based ETFs*, which are ETFs whose underlying securities are stocks that compose a market index. A *market index* is a value meant to track some aspect of the stock market. Some examples of market indexes are the Standard and Poor's (S&P) 500 and Dow Jones Industrial Average (DJIA) which track market performance, or the Volatility Index (VIX) which tracks market volatility. An index-based ETF is designed to track its corresponding market index, but an ETF also has a trading price, because it trades as a symbol on a stock market. When the market is open, a market participant can observe the trading price of the ETF and calculate the current value of the

33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada.

market index through the underlying symbols, which in turn creates arbitrage opportunities between the ETF's trading price and the market index.

Arbitrage opportunities can help an ETF's trading price track its corresponding market index [Antoniewicz and Heinrichs, 2014], but they also tether the price volatility of an ETF's underlying symbols [Ben-David et al., 2015]. The rise in market-wide price volatility throughout 2018 and 2019 has become an increasingly scrutinized topic as some posit an imminent recession and a decrease in investor trust [Carlson, 2019, Li, 2019]. This recent increase in market-wide volatility raises questions about the role ETFs may play in amplifying this volatility through price co-movements in their portfolios. We specifically focus on *mini flash crashes*, which are short volatility events where the price rapidly drops, then quickly reverts back to a similar price as preceded the drop.

Hypothesis: We hypothesize that the presence of an ETF in a market serves to transmit mini flash crashes from one underlying symbol to other underlying symbols in its portfolio.

To explore this question we utilize an autonomous agent-based model. Our market consists of multiple stocks which compose the portfolio of an ETF. We study our simulated market with and without an ETF. We introduce a mini flash crash to a single underlying symbol by way of a trader quickly submitting a series of large marketable limit orders to determine if arbitrage strategies can impact the prices of the other underlying symbols.

Prior work examining ties in market volatility to ETFs has mostly used quantitative models and historical data [Ben-David et al., 2015, Da and Shive, 2017, Lynch et al., 2019, Madhavan and Morillo, 2018]. Using a simulated model instead allows us to examine otherwise identical market environments with and without an ETF present. When an ETF is present, we implement a trading strategy that trades on arbitrage opportunities between the ETF and its underlying symbols. This allows us to observe whether trading strategies dependent on an ETF and its portfolio directly impact their underlying symbols' price volatility.

We employ a simulated market model using a standard limit order book and message system similar to the US stock exchange NASDAQ [Byrd et al., 2019]. Our market contains two symbols which compose an ETF's portfolio. We run our model with and without the ETF present in the market. Our model is populated with 100 background agents per symbol, and when an ETF is present there are 50 arbitrageurs which trade only on the ETF and 50 arbitrageurs which trade on both ETF and underlying symbols. We also use one impact agent to submit a series of large, marketable orders to create a mini flash crash.

We find that the presence of ETFs do increase the volatility of their underlying symbols. When one underlying symbol experiences a mini flash crash, the other symbol also experiences a mini flash crash in the opposite direction when an ETF is present. We also observe that the overall volatility of underlying symbols increases. This demonstrates that ETFs can spread volatility events through their portfolios and contribute to market-wide volatility.

This paper is organized as follows. We next discuss prior work in Section 2. In Section 3 we provide an overview of an ETF's market structure and the arbitrage opportunities this structure creates. We then describe the market mechanism for our simulated ETF environment in Section 4. Section 5 presents our findings on the impacts of mini flash crashes in a market with and without an ETF. Lastly, we conclude in Section 6.

2 Related Work

Our market model is based off of and expands on that of Byrd et al. [2019]. We also use background agents that adopt a zero-intelligence (ZI) strategy [Gode and Sunder, 1993, Farmer et al., 2005]. Our exploration of transparency and explainability of ETFs' impact on market volatility is different than the questions addressed in the previous work of similar simulated market models.

The majority of prior work uses historical data and quantitative models to analyze the impact of ETFs on the volatility of their underlying symbols. Ben-David et al. [2015] constructed a quantitative model and used historical data to conclude that ETFs increase the volatility of their underlying symbols. Da and Shive [2017] also analyzed historical data and found that ETFs contribute to price co-movement in the symbols in their portfolios. In contrast, Madhavan and Morillo [2018] found that underlying symbol price co-movements correlate to macro-market movements, rather than the presence of the

ETF. Lastly, Lynch et al. [2019] implemented a real trading strategy and discovered that traders can make a profit using ETF arbitrage because of portfolio price co-movements.

There are also previous studies examining what leads to ETF arbitrage opportunities. Using historical data, Box et al. [2019] found that a price shock or order imbalance in an ETF's underlying portfolio typically precedes ETF arbitrage opportunities. In our study we use mini flash crashes, a type of price shock, to study the market impacts of ETF arbitrage. Marshall et al. [2013] studied historical data and found that the difference between the best visible bid and offer increases before ETF arbitrage opportunities.

To our knowledge, there is little prior work examining the impact of ETFs on market volatility in a simulated environment. We study this question with an autonomous agent-based model because we can then study the exact same market environment with and without an ETF. We can also be certain in a simulated environment that volatility in the underlying does not stem from externalities.

3 ETF Market Structure

The market structure of an ETF is unique in that ETFs actually trade on two markets. An ETF trades on the stock market like any other stocks, but rather than its value being derived from a company, the value is based on a portfolio of other securities. Similar to other portfolio management funds, such as hedge funds, an ETF also allows investors to accrue and liquidate fund holdings. This fund market is referred to as an ETF's primary market, and the stock market is the ETF's secondary market. Since an ETF trades on two different markets, it has two potentially deviating prices. This creates an arbitrage opportunity for market participants able to detect and act on deviations between the ETF's two prices.

3.1 Primary Market

An ETF is a portfolio management fund, and like various types of these funds, a select few participants, commonly referred to as authorized participants (AP), can invest and divest in an ETF at a time and price determined by its fund managers [Novick et al., 2017]. An AP can trade daily at the *Net Asset Value* (NAV), the price of the ETF in the primary market. For an index-based ETF, the NAV is calculated by finding the true value of the market index at the close of the stock market. Once the NAV is determined, the fund will accept orders until a predetermined time when all orders are executed. These orders are referred to as *basket orders*, which are composed of shares of the ETF or shares of the underlying symbols. APs submit basket orders to the fund manager in return for shares of underlying symbols or shares of the ETF, respectively [Gastineau, 2004]. To submit a basket order, an AP must first acquire shares of the ETF or underlying symbols on the stock market. After a basket order is executed, the AP can trade its newly obtained shares on the stock market.

A notable aspect of the ETF primary market is that like other portfolio management funds there are high barriers to entry. APs tend to be large institutional investors because they need to have a lot of capital and resources to create basket orders on an ETF's primary market [Antoniewicz and Heinrichs, 2015]. While the primary market is fairly exclusive, ETFs are unique for portfolio management funds in that their secondary markets are much more accessible to a variety of investors.

3.2 Secondary Market

In contrast to other portfolio management funds, ETFs also trade on a secondary market, which is the stock market. A stock share of an ETF corresponds to a small portion of the ETF's portfolio. On the stock market, an ETF trades like any other symbol where traders can submit orders to buy or sell shares of the ETF on a stock exchange [Engle and Sarkar, 2006]. When a trader submits an order, it will either transact with an order on the opposite side, or will rest in a limit order book until it can be transacted or the trader cancels the order. Since an ETF has an active order book throughout a trading day, it also has a trading price on its secondary market. We define the *trading price* of an ETF as the midpoint between the highest bid price and lowest offer price for visible volume in its secondary market's limit order book.

An ETF's secondary market gives an ETF much lower barriers to entry than other portfolio management funds [Poterba and Shoven, 2002]. Any investor who can gain access to the stock market, which is typically done through a broker, can invest in an ETF. The accessibility of the secondary market is a principal reason for the popularity of ETFs [Golub et al., 2013]. Participants in the primary market can also trade in the secondary market.

3.3 Arbitrage between the Primary and Secondary Markets

An ETF trades at two prices, the NAV and trading price on its primary and secondary markets, respectively. Even though the NAV is only calculated once per day, it can be estimated throughout the secondary market's trading day by calculating the market index. When the trading price drifts from the market index, a trader can make a profit by buying the lower priced asset and selling the higher priced asset. Arbitrage between these two prices is encouraged because it helps force the trading price to more closely track the index [Antoniewicz and Heinrichs, 2014].

While any secondary market participant can arbitrage price deviations between the market index and trading price, this is a particularly beneficial arbitrage opportunity for participants with access to both the primary and secondary market [Poterba and Shoven, 2002]. Such a participant can arbitrage in the secondary market, then submit basket orders in the asset it is long in for shares of the asset it shorted. It can cover its shorts in the secondary market, making it overall net zero.

4 Market Mechanism

4.1 ABIDES

To investigate the potential for an ETF to spread a mini flash crash from one symbol to another, we employ a discrete event market simulation constructed on the ABIDES platform [Byrd et al., 2019]. [1] The platform provides a continuous double auction market with securities priced in cents, a set of typical background agents, and a kernel which drives the simulation with nanosecond resolution while permitting sparse activity patterns to be efficiently computed.

For the ETF secondary market, we use the provided ABIDES exchange agent, which operates in a manner similar to the NASDAQ. The market is open from 09:30 to 16:00, lists any number of securities for trade, and provides a distinct order book mechanism for each security. The exchange accepts limit orders of any share volume, and cancellation of same, and transacts (including partial execution) those orders against a security's limit order book with a typical price-then-FIFO matching algorithm. The exchange responds to requests for market hours, last trade prices, and market depth quote requests, with depth one representing the current best bid and ask.

4.2 Primary Market in ABIDES

After the exchange agent stops accepting orders, the primary market receives the close price of each underlying symbol in its ETF's portfolio. The primary market uses these closing prices to calculate the value of the index and uses this value as the NAV. Then the primary market opens for basket orders. Every basket order it receives is executed at the NAV, and the agent is notified of its transaction.

4.3 Market Index

We define this model's market index at time t as the sum of the price of a bundle of stocks:

$$\iota_t = \sum_{i=1}^{n} p_{i,t}. \tag{1}$$

We based our market index off the DJIA, which is a US stock market index that tracks 30 large market capitalization stocks. However, our index is simpler than the DJIA, which sums the prices of these 30 stocks and divides by the DJIA divisor. Our market index calculation simplifies the basket orders in the ETF primary market, because one share of the ETF is equivalent to a basket of one share of each symbol in its portfolio.

[1] The ABIDES source code is available at `https://github.com/abides-sim/abides`.

4.4 Exogenous Price Time Series

Similar to Chakraborty and Kearns [2011], we assume the existence of a mean reverting exogenous price time series or *fundamental value series* for each security, which represents an immeasurable consensus of the "true" value at any point in time. In this work, we utilize a "sparse discrete" fundamental which follows an Ornstein-Uhlenbeck (OU) process [Uhlenbeck and Ornstein, 1930] augmented with periodic "megashocks" of higher variance. Given a prior value for the sparse discrete fundamental, a value can be obtained for any subsequent time t by sampling from a normal distribution with:

$$\mathbf{E}[Q_t] = \mu + (Q_0 - \mu)e^{-\gamma t} \tag{2}$$

$$\mathbf{Var}[Q_t] = \frac{\sigma^2}{2\gamma}(1 - e^{-2\gamma t}) \tag{3}$$

as shown in Chakraborty and Kearns [2011], where μ is the long-term mean fundamental value, Q_0 is the prior fundamental value from which we wish to sample a new value for time t, a mean reversion rate γ, and a volatility value σ. This fundamental process has a benefit over the typical discrete mean reverting process when time resolution is substantially finer than the typical agent arrival rate, in that we can obtain Q_t directly from Q_0 without the requirement to compute all intermediate $Q_1 \ldots Q_{t-1}$. That is to say, it permits the simulation to "skip time" whenever no agents will arrive at the market.

The OU process produces a time series with a single scale of variance, which we configure to provide appropriate noise during mean reversion. To obtain a price time series more similar to a stock price over a longer window, we augment the OU process with the application of stochastic "megashocks" as in Byrd [2019] that arrive according to a Poisson process and apply price deviations drawn from a higher variance bimodal normal distribution with mean zero, and positive and negative modes substantially away from zero. This is intended to represent exogenous events which can substantially alter value perception and periodically provides much larger price moves from which the OU process can revert.

4.5 ETF Fundamental Calculation

Most symbols in our market model have independent fundamental values, but an ETF's fundamental is dependant on the value of other symbols in the market. Our ETF is an index-based ETF, and we use the index from Equation 1. Using this index and the value of the fundamental of symbol i, we can find the fundamental of the ETF at time t by:

$$f_t = \sum_{i=1}^{n} Q_{i,t}. \tag{4}$$

4.6 Background Agents

We employ a population of background traders for our experiments drawn from the Zero Intelligence (ZI) family of strategies [Gode and Sunder, 1993]. Our background ZI agents are modeled after those found in Wah et al. [2017]. The ZI traders arrive according to a Poisson process with λ_a, cancel any outstanding orders, and place an order with equal probability to buy or sell. Each ZI trader is assigned a single security of interest. When arriving at the market, an agent receives a noisy observation of the exogenous price time series described above and uses this observation to update an individual Bayesian belief concerning the current fundamental value of the security. The agent then projects this belief forward to the end of the market trading period to obtain an estimated final valuation for the security. An agent uses this final valuation and a private valuation to select a limit price.

To ensure that an order will be gainful, the ZI agent randomly selects a level of surplus $R \in [R_{\min}, R_{\max}]$ to demand, and shades its bid to produce that surplus if transacted, or zero if not transacted. Each agent also has a specified hyperparameter η. If a market order would guarantee surplus of at least ηR, the agent will place that order instead of its computed limit order.

4.7 Impact Agent

We implement an impact agent to induce a mini flash crash. Beginning at time τ, this agent submits a rapid series of n marketable sell orders of size q, δ seconds apart. This method consumes the bid side of the order book, causing the price to rapidly drop. Once other trades continue to submit orders, the price typically recovers.

4.8 ETF Arbitrage Agent

We develop two arbitrage strategies between the ETF and its underlying symbols. Similar to the background agents, these arbitrage agents arrive to the market according the a Poisson distribution with λ_a. When an arbitrage agent enters the market at time t, it calculates the difference between the ETF trading price, p_t, and the market index ι_t:

$$\Delta_t = p_t - \iota_t. \tag{5}$$

These agents only submit marketable orders, so if they sell an asset they will price it at the best available bid, and if they buy they will price their order at the best available offer.

ETF Only Arbitrage This agent will only trade on the ETF when it arbitrages, and it makes trading decisions with some $\varepsilon \geq 0$:

$$\begin{cases} \Delta_t > \varepsilon & \text{Sell ETF,} \\ \Delta_t < -\varepsilon & \text{Buy ETF.} \end{cases}$$

ETF Market Maker The market making agent trades both the ETF and its underlying symbols. Similarly to the ETF only arbitrage agent, it makes trading decisions with some $\varepsilon \geq 0$:

$$\begin{cases} \Delta_t > \varepsilon & \text{Sell ETF and buy underlying symbols,} \\ \Delta_t < -\varepsilon & \text{Buy ETF and sell underlying symbols.} \end{cases}$$

These agents trade on the primary market. When the primary market opens, they receive the NAV, then decide if they should submit basket orders by comparing the NAV and close price of the ETF, p_T, where T is the time of the secondary market close. More formally:

$$\begin{cases} NAV - p_T > 0 & \text{Basket order of ETF shares for underlying symbol shares,} \\ p_T - NAV > 0 & \text{Basket order of underlying symbol shares for ETF shares.} \end{cases}$$

If they submit basket orders, then they can hopefully end the day net zero. The market makers need to be the fastest agents in the system in order to be profitable when trading on so many symbols.

5 Experimental Results

5.1 Market Environment Settings

We explore two market environments, both have one exchange agent and two symbols, but one environment also contains an ETF whose portfolio is composed of the two other symbols. The market environment which contains the ETF, also has an ETF primary market. These two market environments are identical, except one contains an ETF while the other does not. The exchange agent accepts orders between 09:30 and 16:00. The ETF primary market accepts orders between 17:00 and 17:01.

Given the nanosecond resolution of ABIDES, we use the sparse fundamental previously discussed. For each symbol which is not an ETF, we produce its global fundamental from Equations 2 and 3,, with fundamental mean $\mu = 10^5$, mean reversion $\gamma = 1.67 \times 10^{-13}$, and market shock $\sigma = 0$. The sparse fundamental experiences a series of megashocks throughout the trading period, and these arrive according to a Poisson distribution with $\lambda = 2.77778 \times 10^{-13}$. We draw the size of these megashocks from a binomial normal distribution with means $\mu_{s,1} = 0$ and $\mu_{s,2} = 10^3$, and variance $\sigma_s = 5 \times 10^4$. Each background agent can observe each fundamental with observation noise $\sigma_n = 10^6$.

Each symbol in this market is populated with 100 background agents, so in a market environment with an ETF there are a total of 300 background agents and an environment with no ETF there are

Table 1: Strategies employed by the background traders.

Strategy	ZI_1	ZI_2	ZI_3	ZI_4	ZI_5	ZI_6	ZI_7
$R_{\min}$	0	0	0	0	0	250	250
$R_{\max}$	250	500	1000	1000	2000	500	500
η	1.0	1.0	0.8	1.0	0.8	0.8	1.0

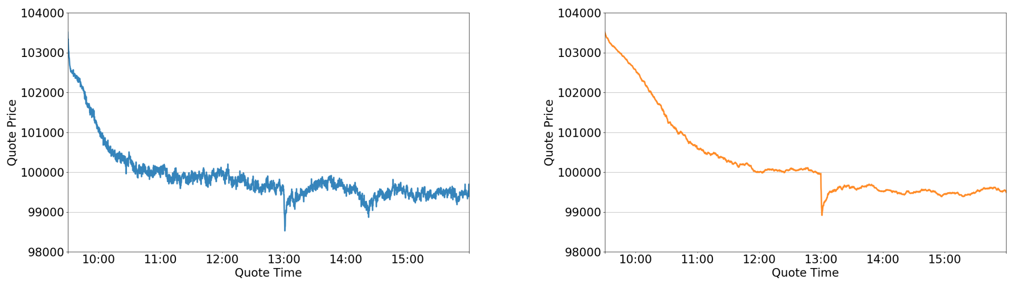

(a) Market contains an ETF and arbitrage agents. (b) Market does not contain ETF and arbitrage agents.

Figure 1: Average price time series over 100 simulation runs of underlying symbol s_0. This symbol experiences a mini flash crash at 13:00 through an impact agent submitting a series of large trades.

200 background agents. Table 1 specifies the strategies of the background agents. The background agents arrive to the market according to a Poisson distribution with $\lambda_a = 10^{-12}$. ABIDES lets us control the message delays of each agent, and the background traders have the highest latency in the market. These agents submit orders of size $q \in [100, 200, \ldots, 1100]$, but can hold a maximum number of units at any time $q_{\max} = 10^3$. The size of each order is drawn from a uniform random distribution. Lastly, the private value variance is $\sigma_{PV} = 5 \times 10^6$.

We create a mini flash crash in one underlying symbol with a single impact agent, which submits its first trade at $\tau =$ 13:00. It then submit $n = 10$ trades with size $q = 600$ and $\delta = 6$ seconds between each trade. This agent has low latency so there is little information delay between different market instances.

We implement two ETF arbitrage strategies when there is an ETF in the market, with 50 ETF only arbitrage agents and 50 ETF market maker agents. For both of these strategies $\varepsilon = 250$, and they submit orders of size $q = 100$. The ETF only arbitrage agents arrive to the market according to a Poisson process with $\lambda_a = 5 \times 10^{-3}$, while the ETF market makers arrive with $\lambda_a = 5 \times 10^{-5}$. Both of these traders have lower latency than background traders, but the market makers are faster and have the lowest latency in our market environment.

5.2 Results

In this paper we analyze ETFs' impact on their underlying symbols when one symbol experiences a mini flash crash. To assess this impact we examine and compare the price of the underlying symbols in environments with and without an ETF. We take the average price of each underlying symbol over 100 simulation runs, which allows us to make more concrete conclusions on the impact of the ETF on its underlying symbols' volatility.

Figure 1 depicts the average price of s_0, which is the underlying where an impact agent enters the market at 13:00 and creates a mini flash crash. Figure 1b shows s_0 when there is an ETF and ETF arbitrage strategies in the market, while Figure 1b shows when there is not an ETF in the market. In both market environments it is clear the impact agent causes a distinct mini flash crash in s_0. However, s_0's price when arbitrage agents are present is more volatile throughout the day. This is likely a result of the high number of ETF market makers in the market, and that these agents move the order book more than the background agents because they only submit marketable orders.

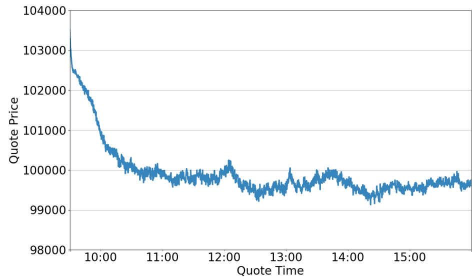

(a) Market contains an ETF and ETF arbitrage agents.

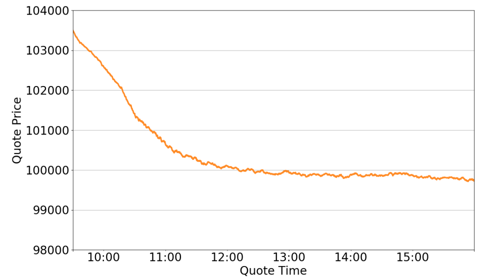

(b) Market does not contain ETF and arbitrage agents.

Figure 2: Average price time series over 100 simulation runs of underlying symbol s_1. An impact agent does not trade on this symbol.

Figure 2 depicts the average price of s_1, which is the symbol in the ETF's portfolio that does not possess impact agent activity. Around 13:00 when the impact agent trades on symbol s_0, Figure 2b shows that symbol s_1 sees no price volatility or reaction to the mini flash crash in s_0. However, when there is an ETF and ETF market makers present, Figure 2a shows that symbol s_1 reacts to the mini flash crash by experiencing its own volatility event in the opposite direction. At 13:00 s_1 momentarily increases in price at a smaller magnitude than the price drop in s_0. This demonstrates that while an ETF exists, the price of the underlying symbols do not necessarily move in the same direction, but they do move together. Therefore, ETFs do impact and increase the market volatility in their underlying symbols by spreading events like mini flash crashes through their portfolios.

6 Conclusion

We analyze a simulated market model with a stock market and ETF primary market. We explore two market environments, one with an ETF and one without an ETF, though both environments contain two symbols which compose the ETF's portfolio. The market is also populated with numerous background agents, an impact agent which creates a mini flash crash, and ETF arbitrage agents. These arbitrage agents only trade in the environment when an ETF is present. We find that the presence of an ETF in a market can transmit a volatility event, like a mini flash crash, throughout its underlying symbols. Though we do find that the other underlying symbols experience a price change in the opposite direction of the symbol where the mini flash crash originated.

A limitation of this study is that the size of our ETF portfolio is only two symbols. For most ETFs in the real-world, their portfolios are significantly larger, and we might get more insight about real ETFs if we increased the size of our portfolio. Another limitation is the limited type of background agents. We only implement ZI agents, which do not consider the order book or previous transactions, so these background agents do not change their trading decisions based on a mini flash crash. It could be beneficial to also utilize trading strategies dependent on the order book and price movement, because these agents might exacerbate mini flash crashes.

We hope to further this work and explore this question with learning and smarter agents. Currently our arbitrageurs to not consider whether trading is beneficial to their profits later on, and it could be interesting to develop an arbitrageur who learns better times to take advantage of arbitrage opportunities. It could also be interesting to create a background agent who attempts to learn when the price is about to change, so these agents could potentially avoid adverse selection by updating and canceling orders leading into volatility events. Lastly, it would be insightful to find the market equilibrium to determine if certain traders are more profitable than others in a market with an ETF.

Acknowledgements

This material is based on research supported in part by the National Science Foundation under Grant no. 1741026, and by a JPMorgan AI Research Fellowship.

This paper was prepared for information purposes, in part, by the AI Research Group of JPMorgan Chase & Co and its affiliates ("J.P. Morgan"), and is not a product of the Research Department of J.P. Morgan. J.P. Morgan makes no explicit or implied representation and warranty and accepts no liability, for the completeness, accuracy or reliability of information, or the legal, compliance, financial, tax or accounting effects of matters contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction.

References

Rochelle Antoniewicz and Jane Heinrichs. Understanding exchange-traded funds: How ETFs work. *SSRN Electronic Journal*, pages 1–40, 2014.

Rochelle Antoniewicz and Jane Heinrichs. The role and activities of authorized participants of exchange-traded funds, 2015. Investment Company Institute.

Itzhak Ben-David, Francesco Franzoni, and Rabih Moussawi. Do ETFs increase volatility?, 2015. Securities and Exchange Commission Comments.

Tavis Box, Ryan Davis, Richard B. Evans, and Andrew A. Lynch. Intraday arbitrage between ETFs and their underlying portfolios. *SSRN Electronic Journal*, pages 1–73, 2019.

David Byrd. Explaining agent-based financial market simulation. *eprint arXiv:1909.11650*, pages 1–10, 2019.

David Byrd, Maria Hybinette, and Tucker Hybinette Balch. ABIDES: towards high-fidelity market simulation fro AI research. *eprint arXiv:1904.12066*, pages 1–13, 2019.

Ben Carlson. Spike, swoon, repeat: What to make of a volatile stock market—that has basically gone nowhere for a year and a half, 2019. Fortune.

Tanmoy Chakraborty and Michael Kearns. Market making and mean reversion. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 307–314. ACM, 2011.

Zhi Da and Sophie Shive. Exchange traded funds and asset return correlations. *European Financial Management*, pages 136–168, 2017.

Robert Engle and Debojyoti Sarkar. Premiums-discounts and exchange traded funds. *The Journal of Derivatives*, 13:27–45, 2006.

J Doyne Farmer, Paolo Patelli, and Ilija I Zovko. The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences*, 102:2254–2259, 2005.

Gary L Gastineau. The benchmark index ETF performance problem. *The Journal of Portfolio Management*, 30:96–103, 2004.

Dhananjay K Gode and Shyam Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101:119–137, 1993.

Ben Golub, Barbara Novick, Ananth Madhavan, Ira Shapiro, Kristen Walters, and Maurizio Ferconi. Exchange traded products: Overview, benefits, and myths, 2013. Black Rock.

Yun Li. Greenspan says the stock market will determine whether there's a recession, 2019. CNBC News.

Hailey Lynch, Sébastien Page, Robert A. Panariello, James A. Tzitzouris Jr., and David Giroux. The revenge of the stock pickers. *Financial Analysts Journal*, 75:34–43, 2019.

Ananth Madhavan and Daniel Morillo. The impact of flows into exchange-traded funds: Volumes and correlations. *The Journal of Portfolio Management*, 44:96–107, 2018.

Ben Marshall, Nhut H. Nguyen, and Nuttawat Visaltanachoti. ETF arbitrage: Intraday evidence. *Journal of Banking and Finance*, 37:3486–3498, 2013.

Barbara Novick, Ananth Madhavan, Samara Cohen, Sal Samandar, Sander Van Nugteren, and Alexis Rosenblum. A primer on ETF primary trading and the role of authorized participants, 2017. BlackRock.

James Poterba and John Shoven. Exchange traded funds: A new investment option for taxable investors. *American Economic Review*, 92:422–427, 2002.

George E Uhlenbeck and Leonard S Ornstein. On the theory of the brownian motion. *Physical review*, 36(5):823, 1930.

Elaine Wah, Mason Wright, and Michael P Wellman. Welfare effects of market making in continuous double auctions. *Journal of Artificial Intelligence Research*, 59:613–650, 2017.